Open in app ↗ Sign up Sign in

Medium Search

Uniswap + Defi + Ethereum + Quantitative Finance +

A Guide for Choosing Optimal Uniswap V3 LP Positions, Part 2

8 min read · Sep 15, 2021

Guillaume Lambert Follow

Listen Share

This is the second article in a series of medium posts about choosing and managing Uniswap v3 LP positions. Read part 1 here.

TL;DR: Uni v3 liquidity positions have to balance capital efficiency and the time spent in-the-money. The product of these two effects results in a simple √T dependence for narrow liquidity ranges. Holding a position for a long time and/or deploying to wider ranges reduce LP returns, a cost we call ***radical liquidity drag****.*

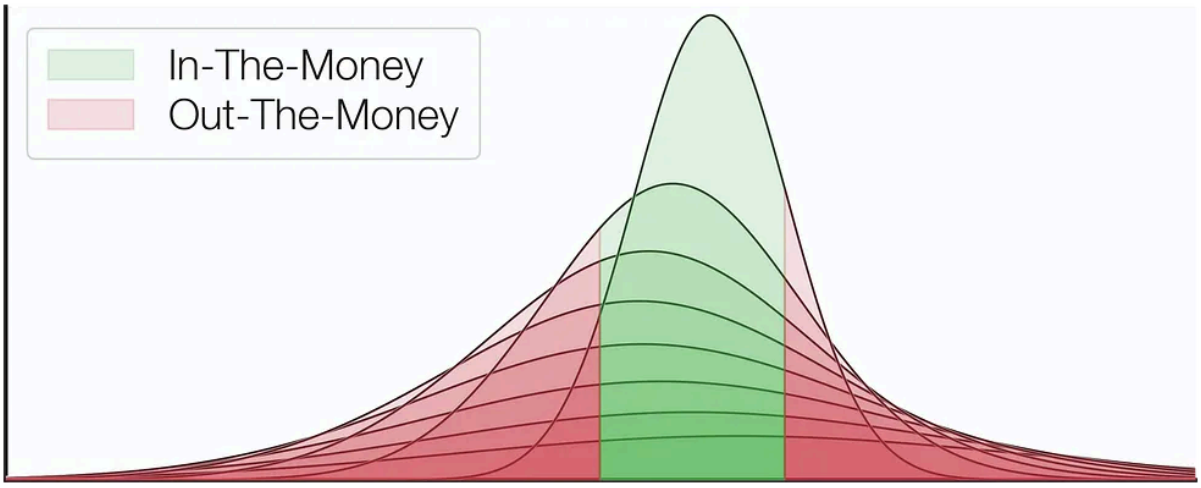

In our previous post, we derived an expression describing the probability of a Uniswap v3 position will end up within the upper/lower tick after a time T. Since fees are continuously accrued when a position is within range of the LP position, perhaps a more sensible question to ask is: ***what fraction of time is spent within a specific range between times 0 and T***?

In this article, we will derive an expression for the expected amount of fees to be accrued by a LP position. We will find that the amount of fee is a combination of the *effective liquidity*, which depends on the upper/lower ticks of the LP position, and the *time spent in-the-money,* which depends on the volatility of the underlying asset.

An overview of Uniswap Fees

Liquidity providers collect fees based on trading volume. The total amount of collected fees during a given time window should depend on

1. The fee tier: 0.05%, 0.3% or 1%
2. The trading volume.
3. The total liquidity at the traded tick.
4. The LP position's liquidity.
5. Fraction of time spent inside the upper/lower ticks.
6. The total amount of time the LP position is deployed

We can separate the expression for the amount of fees collected into three components, which can be written as:

$$\text{LP returns} = \underbrace{\left(\text{fee tier} \cdot \frac{\text{Trading Volume}}{\text{Total Liquidity}}\right)}_{\text{fee accrual rate}} \cdot \underbrace{(\text{Position Liquidity}) \cdot \left(\frac{\text{T}_{\text{itm}}}{\text{T}_{\text{total}}}\right)}_{\text{effective liquidity}} \cdot \underbrace{(\text{T}_{\text{total}})}_{\text{duration}}$$

- **Fee accrual rate:** the fee tier, trading volume, and total liquidity are dictated by the AMMs smart contract and the market. Thus, when deploying liquidity, users should target Uniswap pools that optimize this quantity (ie. consider the volume, TVL and fee tier).
- **Effective liquidity:** The effective size of a position is the product of the amount of assets locked in each tick of the position and the total time spent inside the

range of the liquidity position. This is controlled by the liquidity provider when the LP position is deployed.

- **Duration**: this is simply the amount of time a position is held before it is removed or rebalanced.

Since the fee accrual rate is not under the control of the user, the only way to optimize a LP position is to optimize the effective liquidity and remain in the position for a duration that maximizes returns.

Capital Efficiency and Position Liquidity

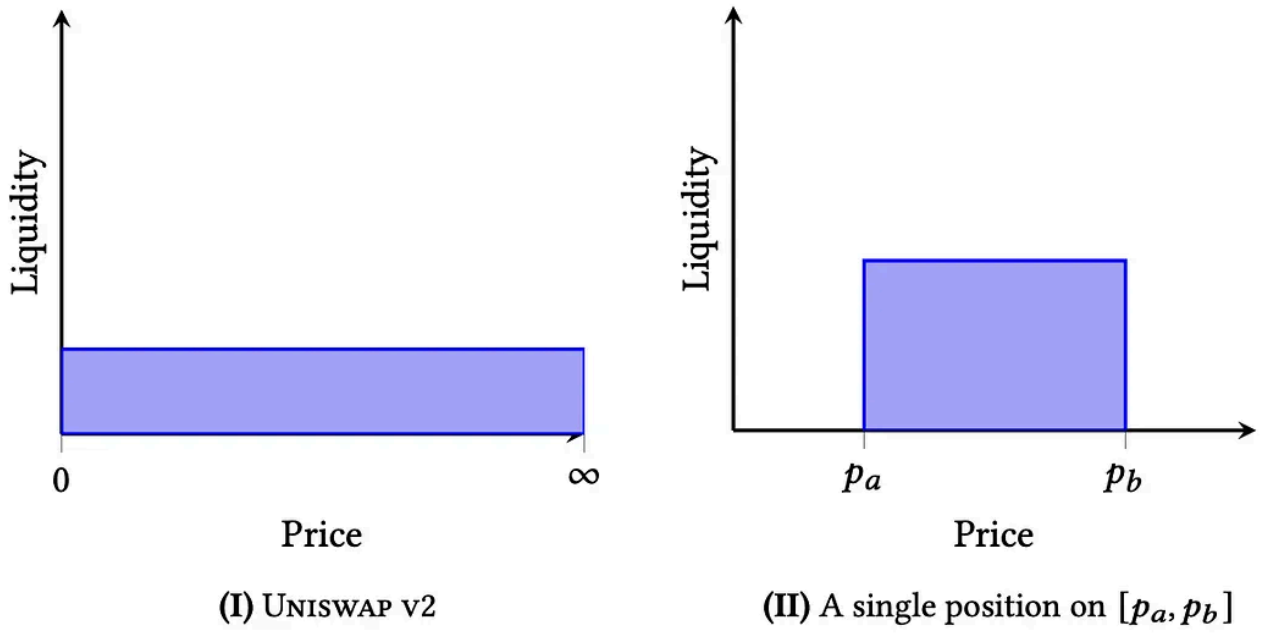

Deploying liquidity in Uniswap v3 is more capital efficient than in Uniswap v2. Source: Uniswap v3 whitepaper.

In Uniswap v2, liquidity is spread across all possible prices, from 0 to ∞, meaning that a large fraction of the liquidity provided may never be "used" because price may not be traded at that specific price. Uniswap v3 set out to fix these capital inefficiency issues by allowing LP to deploy concentrated liquidity between any two "price ticks".

In Uniswap v3, each tick *i* defines a specific price P = 1.0001^i. The index *i* can be any integer number between (-887272, 887272). When deploying a position, a liquidity provider defines a lower tick i_L and an upper tick i_U so that liquidity is deployed between a lower and upper price (t_L, t_H) equal to 1.0001^{i_L} and 1.0001^{i_U}.

Deploying liquidity to a specific range is much more capital efficient than in Uniswap v2, here's Hayden's tweet about this and this analysis by Tienshaoku:

Hence, **the effective liquidity of the position will be much larger for narrow ranges.** Since the total liquidity L is "split" equally into small δL = L/N chunks between all ticks between the upper and lower limits, we can compare the relative capital efficiency of a ranged Uniswap v3 position and a corresponding Uniswap v2 position (the "Full Range" option on the Uniswap v3 interface deploys liquidity between the (-887,272, +887,272) ticks. For a ranged position, the number of liquidity ticks N in a position is given by:

$$\begin{aligned}\text{Number of liquidity ticks } N &= \frac{i_H - i_L}{tickSpacing\ t_s} \\ &= \frac{\ln t_H - \ln t_L}{\ln 1.0001^{t_s}} \\ &= \frac{20,001}{t_s} \cdot \ln r \text{ , where } r = \sqrt{\frac{t_H}{t_L}}\end{aligned}$$

To compare the capital efficiency of Uniswap v3 compared to a "Full Range" position, we plot the number of liquidity ticks as a function of the range factor r for each of the 0.05%, 0.3% and 1% fee tiers to get:

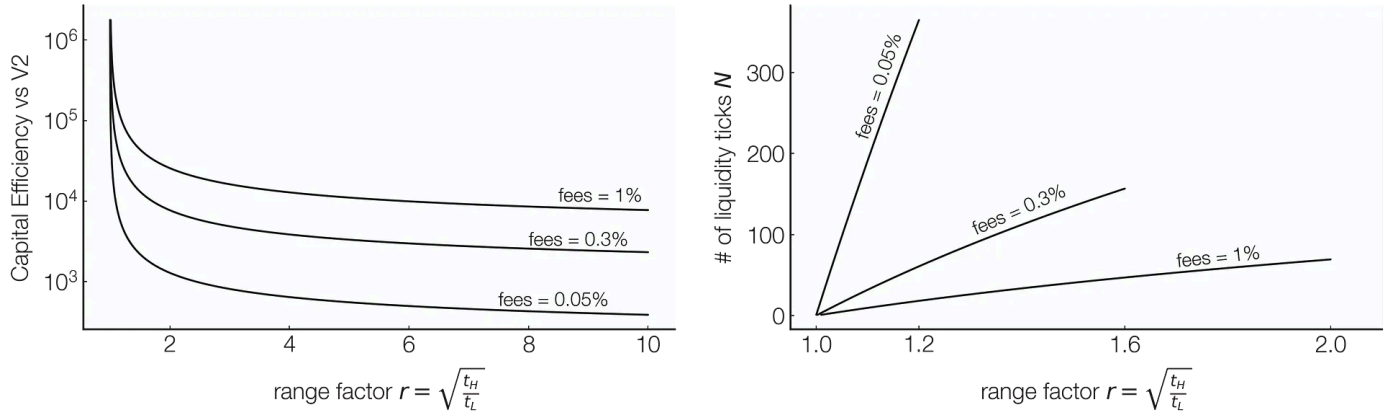

Capital efficiency in Uniswap v3. Left: Deploying to a single tick is >1M times more capital efficient than Uniswap v2. Right: deployed liquidity is equally split into N ticks between the upper and lower limits of the liquidity position.

A couple of takeaways here. First, liquidity deployed to a single-tick in Uniswap v3 is **1,774,544 more capital efficient** than the corresponding Uniswap v2 (ie. Full Range) position. Even a more modest range factor r=1.2 for a 0.3% pool leads to a staggering 29,200-fold improvement in liquidity over the “Full range” option.

Second, extending the range **splits the liquidy** equally among all ticks between the lower and upper ticks. Specifically, the amount δL of liquidity inside each chunk scales as δL~1/ln(r). This means that liquidity deployed using r = 1.2 is effectively 2.5x larger than liquidity deployed using r=1.6.

Get Guillaume Lambert’s stories in your inbox

Join Medium for free to get updates from this writer.

Enter your email

Subscribe

Remember me for faster sign in

Finally, **liquidity fragmentation** can be significant when liquidity is deployed to a smaller fee tier. In other words, deploying liquidity at a range factor r=1.6 results in a splitting of the liquidity into 47 ticks for the 1% pool, 156 ticks for the 0.3% pool, and 940 ticks for the 0.05% pool.

Time spent In-The-Money

The second factor that influences liquidity returns is the time spent inside the deployed liquidity range. While this quantity is typically not relevant in traditional finance (although it is related to the pricing of Asian options), understanding the expected time a position will remain in-the-money is extremely important in Uniswap v3, as fees are only collected when the price is inside the LP range.

We will derive here the expected time spent in the money for a liquidity position. As a reminder, we derived in our previous post the expression for the probability that a price ends within a specific price range defined by a lower tick t_L and an upper tick t_H after a time T is:

$$Prob(\text{In-The-Money}) = \text{Erf}\left(\frac{ln\left(\sqrt{\frac{t_H}{t_L}}\right)}{\sigma\sqrt{2T}}\right)$$

Hence, to find the fraction of the total amount of time spent within that range, we need to integrate the expression for Prob(In-The-Money) over all times and divide by T:

$$\frac{T_{itm}}{T_{tot}} = \frac{\int_0^T Prob(itm)dt}{\int_0^T dt}$$

Solving this integral using an Erf integral table (see Ng & Geller, 1968) or using Mathematica, we obtain, using r = √(tH/tL):

$$\frac{\mathrm{T_{itm}}}{\mathrm{T}} = \frac{\left(\ln^2(r) + \sigma^2\mathrm{T}\right)\text{Erf}\left(\frac{\ln(r)}{\sigma\sqrt{2\mathrm{T}}}\right) + \ln(r)\left(\sqrt{\frac{2}{\pi}}\sigma\sqrt{\mathrm{T}}e^{-\frac{\ln^2(r)}{2\sigma^2\mathrm{T}}} - \ln(r)\right)}{\sigma^2\mathrm{T}}$$

Complicated, I know. We did make a couple of assumptions to get the expression above:

1. We assumed that the starting position is the geometric mean of the upper and lower ticks: S0 = √(tH*tL). The expression will be more complex if S0 is not equal to the range's midpoint
2. The timescale $\sigma^2 T$ is small compared to ln(r). This is the same assumption as in our previous, but once again the T_ITM expression will differ somewhat when that's not the case
3. We assumed that the price drift term μ is equal to 0. This is also a valid assumption since the annual drift is a few percent per year at the most.

While this expression is somewhat hard to understand, it can be easily plotted using mathematica or desmos. See the figure below and this desmos link:

- Limiting assumptions for σ, T, and S0, μ=0: https://www.desmos.com/calculator/xou9ukzymi

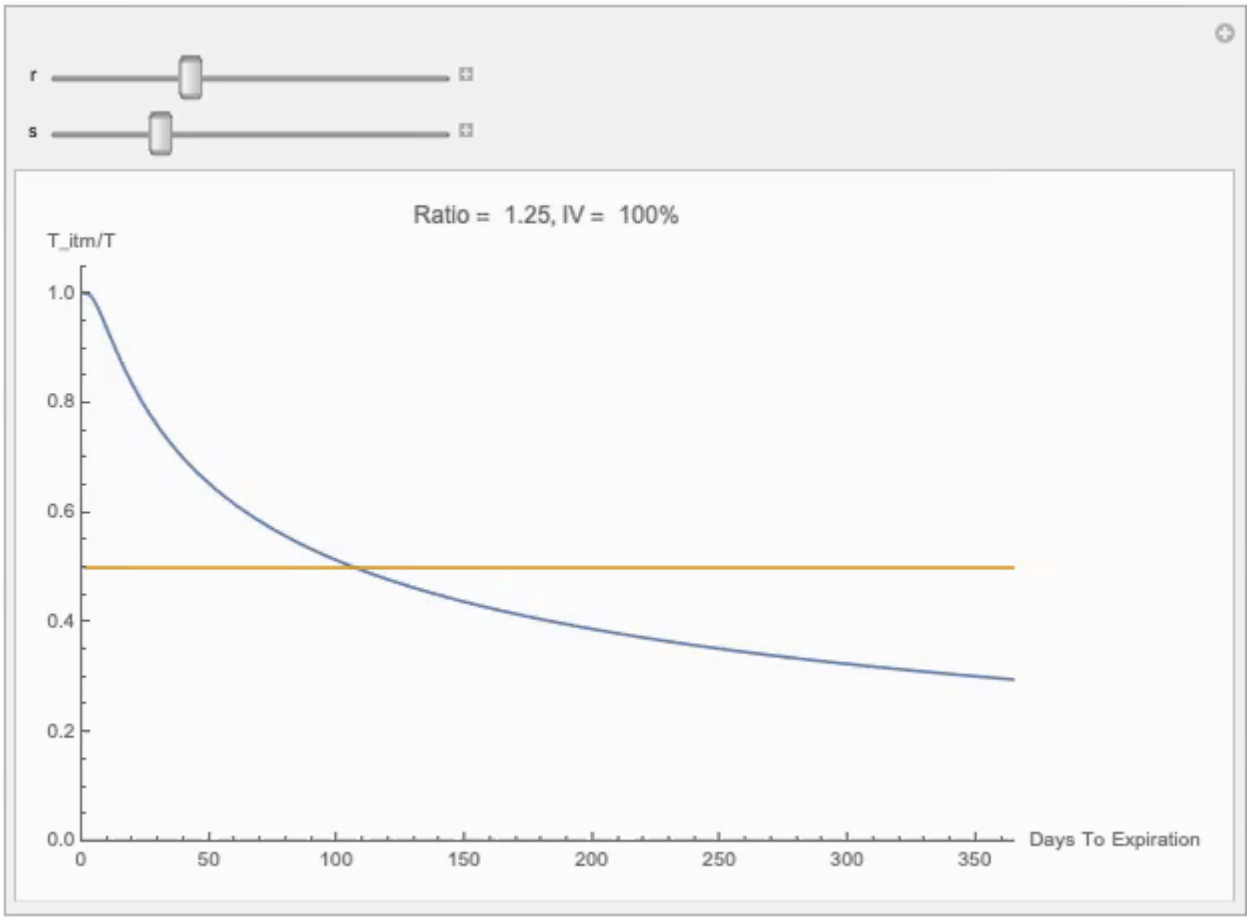

Probability to remain in the money as a function of time, range factor and volatility.

Note that the relative time spent in the money can remain very **close to 100**% for large range factors (ie. r>1.5) and short timescales (<50 days). For r=1.5, the fraction of time spent in-the-money after 1 year is close to 50%. In contrast, if r=1.05, the fraction of time spent in-the-money after 1 year is only a little bit more than 7.5%. Of course, volatility can spike and the value of cryptoassets can double within a few months, rendering this prediction less useful.

Closed form expression for the effective liquidity

We expect the"effective liquidity" to be a combination of both the capital efficiency of the position and the relative amount of time spent in-the-money. So, combining the expression for the liquidity and the time spent in-the-money, we obtain the following expression:

$$\text{LP returns} \quad = \quad (\text{fee accrual rate}) \cdot \left[\frac{L \cdot t_s}{20,001 \ln(r)} \cdot \frac{T_{itm}}{T} \right] \cdot \mathrm{T}$$

Once again, the fee accrual rate is not under the control of the liquidity provider. Only the range factor r and the amount of time a position is maintained is under the control of the user.

We can simplify the expression for the LP returns to get, after some algebra:

$$\text{LP returns} \quad \propto \quad \sqrt{\frac{8}{\pi}} \frac{\sqrt{T}}{\sigma} + \mathcal{O}(r-1)$$

Here's what LP returns look like, graphically (this plot includes all the O(r-1) terms):

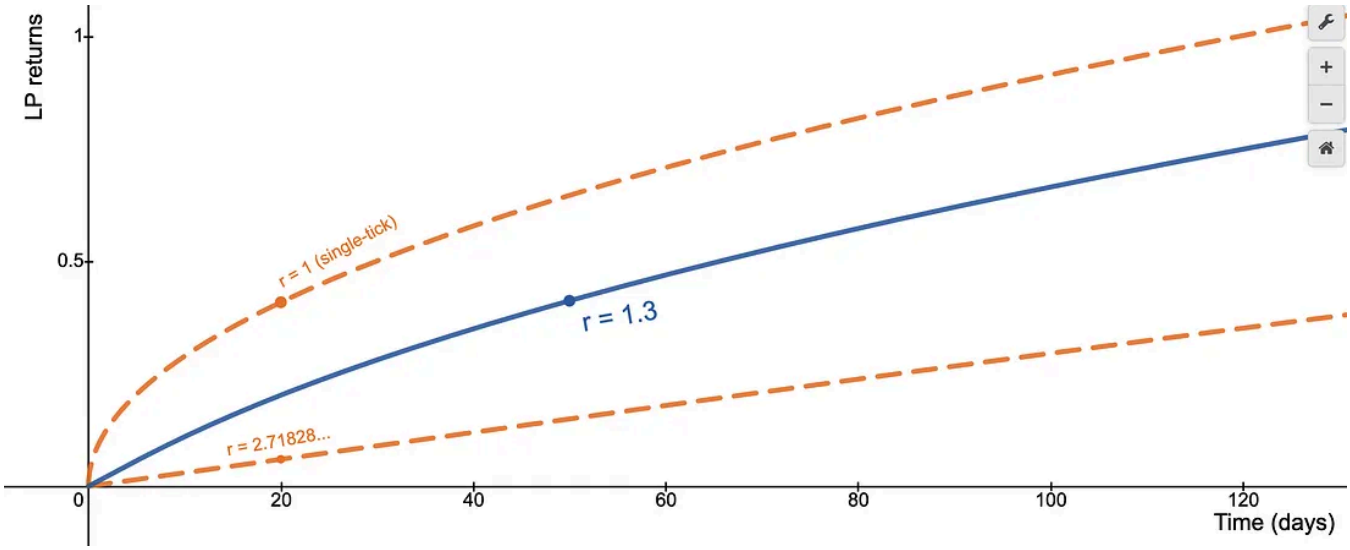

LP returns as a function of the range factor. Liquidity deployed to a single tick results in LP returns with a √T dependence. Returns become linear for larger r, but never exceed those of single-tick liquidity. Desmos: https://www.desmos.com/calculator/bba8lc6vks

These results are a bit surprising for the following reasons:

1. LP returns have a √T dependence. For example, to double the fees collected after 1 day, one needs to wait 3 extra days. We'll call this phenomenon the **radical liquidity drag.**
   Takeaway 1: Exit a position early to avoid radical liquidity drag.

2. LP returns are maximized for very narrow ranges: LP returns for range factors smaller than r~1.05 should lead to a √T dependence that outperforms wider liquidities. Smaller range also helps reduce impermanent loss.
   Takeaway 2: Keep ranges tight.

3. LP returns become linear in time for larger range factors (r>2).
   Takeaway 3:Being lazy and choosing the full range will drastically reduce LP returns.

Conclusion

In this post, we derived an expression for the expected return of a Uniswap v3 LP position. We found that the effective liquidity of a tightly defined liquidity position can be several times more capital efficient than Uniswap v2. We also derived an explicit expression for the time spent inside a specific range.

Interestingly, we found that LP returns for narrowly defined LP positions accumulate fees proportionally to the square root of the time the position is held. This means that a LP position will see diminishing returns over time, and it may be best to exit positions early than keep them on for an extended period of time. This is

perhaps the most interesting point of this analysis: exit your positions early to minimize **radical liquidity drag.**

Another unexpected result is **that fees are maximized for LP positions deployed to a single-tick.** It is important to note that impermanent loss will also be much lower for tight range that are rebalanced frequently compared with wider ranges.

To summarize, a Uni v3 liquidity position has to balance capital efficiency and the time spent in-the-money, and it appears that the product of these two effects results in a simple √T dependence for small r. Larger r, on the other hand, leads to higher order corrections that reduce the LP returns.

Uniswap + Defi + Ethereum + Quantitative Finance +

Follow

Written by Guillaume Lambert

1.7K followers · 22 following

Asst professor in Applied Physics at Cornell. LambertLab.io PI and proud father. Interests: Biophysics, Math, Crypto, and Options (often all at the same time).